



Simplifying Storage Management with Software-Defined Storage

Analyst: David Reine

Management Summary

With all of this snow on the ground in New England, now is the time to think about getting away. It should be a rather simple exercise: contact your local travel agent, if you want, or log onto one of the many travel sites and book an airplane seat, a hotel, and a rental car yourself, with the click of a mouse. If you are considering a business trip, however, travel arrangements can become a bit more complicated. Let's say you are planning to visit five cities in different countries, or even states, in six days. You, the virtual travel agent, may need to make travel arrangements on five different airlines, or more, depending upon connections, forcing you to visit multiple websites. You will need to schedule hotels in each city, perhaps different chains, meaning more, different websites. The same is true for car rentals. Perhaps, you may even need to make train reservations, possibly requiring an even more complex itinerary. Wouldn't it be nice if you could manage even this complex trip within a single website visit?

This same problem confronts the storage administrator in every data center staff of every business: Much like the travel agent, how does he/she manage a heterogeneous mix of data storages devices from a single pane of glass, simply and efficiently. *How can the management of different platforms, using different architectures, be converged into a coherent entity? How can applications disregard the physical infrastructure of each platform, abstracting the data in some logical manner, virtualizing the data into a logical arrangement of critical information?* Furthermore, how can the administrator do this in an intelligent manner, removing the complexities of human intervention? *How can a simplified, cross-platform, virtualized storage solution eliminate the complexity and enable the data center staff at service providers, mid-sized businesses, and enterprises to turn present-day storage chaos into something more rational, both technically and economically?*

The data center staff needs to have the power to seamlessly migrate, recover, protect, and deduplicate data both in the data center and in the cloud, while also being able to communicate with a variety of storage targets, or for that matter, specific hardware, networks, or protocols. In other words, the enterprise data center needs to be able to select the most appropriate storage infrastructure to meet the mission- and business-critical application requirements – both operational and budgetary. The data center must have the capability to meet the business needs while remaining within budget, in order to focus on driving additional business and profitability. The data center staff needs to do more with less (people, time and money), thereby simplifying operations to free up human resources to do things with a higher payback than mundane but necessary storage administration tasks. In addition, the storage infrastructure must facilitate a simplified operation for enterprise partners and customers, dictated internally by the needs of business and not by storage platform vendors. In order to accomplish this, the data center staff must move forward with a more modern approach to storage technology and put their legacy infrastructure in the rear-view mirror. To learn more about what is needed to better meet today's requirements, please read on.

IN THIS ISSUE

> Today's Overwhelmed Data Center.....	2
> Six Storage Challenges.....	2
> Today's Storage Requirements.....	3
> The Solution – Software-Defined Storage.....	4
> Conclusion.....	5

Today's Overwhelmed Data Center

What does the storage in your data center look like today? Let me guess. First of all, you may have a confusing infrastructure made up by possibly, many) storage vendors. Often, this is a result of the isolation of servers to hosting one or a few applications and a concurrent, less-than-ideal (piecemeal), storage strategy. The result is a costly and inefficient storage architecture that, more likely than not, is incompatible with where you need to be. In addition, the deployment of these varied applications and data siloes presents a plethora of on-going, multi-dimensional challenges for the data center staff in administering legacy storage infrastructure. **Yesterday's storage architecture(s) may well be cost-ineffective for today's data center.**

Today's typical data center typically has too many independently-conceived and narrowly-defined storage solutions (often with its own internal or network storage), which has resulted in an architecture that is far more complex than it should be. Furthermore, the deployment of storage platforms from multiple vendors has necessitated the deployment of multiple storage management interfaces, perhaps even causing the need for special-purpose storage administrators, each with unique, platform-specific skills.

Additionally, some of your mission- and business-critical applications have been deployed within a physical environment, while others have been deployed with a more modern virtual environment. **Unfortunately, today's – and tomorrow's – technology doesn't always integrate well with your hodge-podge of legacy storage solutions.**

As a result, the storage options that you now have probably are inadequate for the critical data that needs to be shared throughout your business, including with customers and partners. Beyond your data center, the data in your storage architecture must be shared with other remote sites, probably including the cloud. In fact, the data required by your infrastructure also must be capable of transitioning between internal and external clouds, none of which have to be on-site. The many different storage services that make up your storage infrastructure may not work well across diverse server platforms. A siloed infrastructure does not flow very well across both older and newer technologies in a cross-platform fashion. As a result, your current architecture may well be inadequate for today's storage requirements.

Six Storage Challenges

Because of high administrative and maintenance costs, legacy storage products probably carry a TCO that likely will break the IT budget under today's Big Data demands. There are many challenges facing the data center staff; six will be discussed.

(1) Challenge – The Need for More Storage Capacity

The volume of storage that is required and the urgency to access the stored data has, and always will have, an immediate effect on the TCO of that storage. The issue now is that you have, and will need, much more data, probably with storage requirements doubling every 12-to-18 months. In addition, there is a pressing need to access much of it more quickly than before.

(2) Challenge – The Need for Speedier Storage

These days, information access and decision-making both must be done in a very short period of time, compared to just a couple of years ago. This is why storage assumptions made even last year are causing serious administrative and operational problems today. Now, every enterprise must implement a multi-tiered storage architecture in order to achieve faster response for some applications, while trying to control the costs for less urgent long-term storage.

(3) Challenge – Storage Costs are Rising

Despite the fact that the cost for a GB of storage on spinning media continues to go down on a unit-cost basis, the total storage budget probably is rising because (a) more snapshots and copies are needed, (b) the growth of data is vastly outstripping the unit cost savings, and (c) the demand for more performant storage just costs more. In addition, urgency (getting data in a meaningfully-quick manner) is an important cost multiplier.

(4) Challenge – Data Stored in the Cloud

The advent of maintaining at least some of your data in the cloud presents an entirely new set of challenges, in terms of the security advantages of a private cloud versus a public cloud and also in terms of time to retrieval. Often, the seeming lower costs for using a public cloud may be offset by the increased costs for WAN. As the headlines indicate, there are additional challenges for the data center based upon the reliability of the cloud provider and the length of time required to restore service when the inevitable outage occurs.

(5) Challenge – Multiple Storage Vendors

Mission- and business-critical solutions, that may be incompatible with each other, carry unpredictable costs from vendors who know that they have you locked-in to their storage solutions. Storage platforms from multiple vendors can result in:

- High maintenance costs when the warranty expires;
- Continuing operating costs for storage management software; and, frequently
- On-going headaches for storage administrators charged with managing each of the different storage silos with different administrative tools.

Focusing only on current demands (i.e., “fire-fighting” the hot spots) only makes it worse when you have to upgrade to meet future requirements for more capacity and better performance (largely to handle Big Data needs), and to get rid of costly/aging technology. Multiple vendors surely increase the complexity of IT operations and, as we all know, **complexity is a heavy burden on the data center IT budget.**

(6) Challenge – Self-Interested Storage Vendors

A multitude of storage solutions often put the business at the mercy of storage vendors, who control the release of new technology and often require wholesale replacement to get to the new technology to meet new and near-future needs.

Today’s Storage Requirements

When contemplating the best way to meet these challenges, and the goals of the business, and to contribute to increased profitability, the data center staff must focus on what needs to be done to remove storage complexity and lower the TCO of the IT environment. Emphasis needs to be placed on improving business efficiency through higher performing storage for immediate response and lower cost for higher capacity storage. This boils down to five requirements.

(1) Requirement – Tiers of Storage without Complexity

Nothing is as constant as change. This especially is true in the data center where changes in information technology are occurring on a regular, if not daily, basis. The topology of storage architectures is in a state of flux. Data is (or should be) migrating from one tier of storage to another, from one location to another, as its value and/or demand change. The creation, retention, protection, and processing of data is now critical

to the success of the business. Since we cannot change the business mission or manufacture more time, we need to look to simplify the way we do business in order to lower the TCO. Thus, the IT staff must remove as much complexity from the storage infrastructure as possible.

Data centers around the globe have been moving to tiered storage for quite some time. Solid State Devices (SSDs) have been positioned as Tier-0 storage because of the extremely rapid access that they provide to the most mission-critical data. High-speed rotating disks have been characterized as Tier-1 storage, usually encompassing Fibre Channel and SAS drives, with slower SATA drives typically positioned as Tier-2 devices. For the purpose of this discussion, we will characterize other media, such as tape, as a Tier-3 media. The external cloud, dependent upon WAN access, could be classified as a Tier-4 storage resource, because of the time it takes to move the data to and from a local server.

Because of the significant difference in the cost of these storage tiers, the storage administrator must ensure that data is stored in the appropriate tier, whether based upon the frequency of use or other value-oriented criteria. This requires data movement up and down the tiers, as the demand for it changes, thus aligning the cost of storage to the value of the data and the acceptable time to gain access to it. To do this manually on a dynamic basis is too great and unwieldy a task to even consider.

What type of data, you may ask, requires Tier-0 storage? Obviously, the answer will differ depending upon business requirements; however, one thing is certain, if your mission-critical applications demand *immediate* accessibility, then Tier-0 SSDs (or similar memory-based technologies) are required. When the demand for data on these SSDs becomes *less* immediate, then the migration to disks or tape, or perhaps the cloud, becomes an economic imperative, if you are going to keep overall storage costs under control. The data center staff needs to keep in mind that data time-to-accessibility affects data value and determines the tier of storage required.

(2) Requirement – Data Migration

It should make no difference if you are moving information from a single array to another within the same data center, or representing multiple physical arrays as one or more logical array, including scale-out storage or other storage connected to the server. The data center is looking for a single, simplified solution to handle legacy

storage devices, contemporary devices, and even virtual arrays. The storage migration utility needs to be adaptive in terms of the capability to handle a heterogeneous mix of storage devices, protocols, and operating environments. It also must be network-efficient in order to optimize WAN utilization and space-efficient in terms of storage replication.

(3) Requirement – Business Continuity

The storage solution must be able to provide five ‘9s’ access for application and business availability, even if the underlying devices cannot support this reliability themselves. It must be a single, economical solution for the management and support of all storage. In order to provide always-on business continuity, the solution needs to be able to provide high availability and fail-over for a single-node storage architecture, or a multi-node storage cluster. Furthermore, it needs to be versatile, capable of supporting local, active-active¹, or multi-site environments. A legacy storage architecture (or even several of them) usually cannot always support all of the demands of a modern business. Today’s – and tomorrow’s – solutions must be redesigned to support whatever the future throws at it.

(4) Requirement – Storage Preservation

The protection and recovery of the valuable assets being managed by the data center is vital to the success of every business, regardless of size. No matter where your primary storage resides, storage recovery must be mobile enough to support instant recovery when demanded by mission- and business-critical applications, such as transaction processing, analytics like Hadoop, and data management, and flexible enough to support low-cost recovery for less-demanding requirements. This implies the capability to manage and administer everything from mission-critical data being maintained on Tier-0 SSD devices to more mundane data being backed-up, or archived, to less-speedy tiers, such as tape, or even to the cloud. The data center staff needs the flexibility necessary to provide the custom recovery of systems, applications, and data services. In addition, the administrators need to have a product that is integrated to support physical, virtual, or mixed storage, and capable of being

moved dynamically, whether in-house or in the cloud.

(5) Requirement – Storage Optimization

In order to lower the TCO of the storage infrastructure, the data center staff must ensure that the value of an organization’s data is maximized relative to the tier of storage on which it exists. It is critical for the data center to be able to archive inactive or less urgent (cold) data from more expensive media to tape or to the cloud. It is equally important for the data center staff to be able to migrate this data back to more speedy media when the demand for it becomes hotter. It is also important for the data center staff to be able to improve backup efficiencies by reducing the amount of storage required to be backed-up on a regular basis through data deduplication, as well as archiving. It goes without saying that all of this data migration should be automated to relieve this burden from the data center staff.

The Solution – Software-Defined Storage

Enterprises, mid-sized businesses, and managed service providers (MSPs) all are looking for the same thing – simplified storage management and delivery services that work cross-platform, no matter whose storage devices are being managed or where they reside. These services must be simple and cost-effective, working seamlessly across most storage platforms and data types, no matter modern or legacy, and must have predictable pricing with no “gotchas” when you expand storage capacity or add new functionality.

With the five requirements just discussed as a backdrop, *software-defined storage (SDS)*, when done well, is the most-promising vehicle to help the data center staff to satisfy – simply – the challenges across many diverse storage platforms. SDS is an approach to data storage in which storage management middleware controls storage-related tasks, decoupling storage access and storage management from the underlying physical storage platforms. It should enable the data center to aggregate segregated storage into one shared, global storage resource pool, available to all authorized applications, whether on-site or in the cloud, facilitating the optimization of workloads, performance enhancements, and cost reductions. Because SDS eliminates the constraints of the physical storage platforms, every storage resource can be used more efficiently and its administration can be simplified with unified, automated, policy-based management tools. This result will look like a single storage platform with

¹ A network of independent nodes where each node has access to a replicated database giving each node access. All requests are balanced across the network, as all of the nodes are active (doing meaningful work). Where a failure occurs on a node, another node in the network takes its place, maybe with a reduction in overall performance.

a single view for common storage management, which should free the data center staff for more productive tasks. The pricing for SDS also should be simple, allowing you to be able to predict what the inevitable future capacity growth will cost (without surprises).

Conclusion

Software-defined storage may be the best way to improve the efficiency of enterprise storage, especially when you have a collection of heterogeneous storage solutions already in place. It is all about simplicity via a common, shared approach to all of the storage hardware that it manages.

If you are overwhelmed with paying and caring for storage and share the challenges and requirements presented and discussed, you should look at software-defined storage. Otherwise, prepare for much more of the same.



About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-third year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- **The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.**

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- **Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)**

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, The Clipper Group Calculator, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper’s web pages, posting of Clipper documents on other’s websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts’ travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper’s analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.